arXiv:1305.0218v1 [cs.CV] 1 May 2013

Video Segmentation via Diffusion Bases

Dina Dushnik[1] Alon Schclar[2*] Amir Averbuch[1]

[1] School of Computer Science,

Tel Aviv University, Tel Aviv 69978, Israel

[2] School of Computer Science,

Academic College of Tel-Aviv Yafo, Tel Aviv 61083, Israel

Abstract

Identifying moving objects in a video sequence, which is produced by a static camera, is a fundamental and critical task in many computer-vision applications. A common approach performs background subtraction, which identifies moving objects as the portion of a video frame that differs significantly from a background model. A good background subtraction algorithm has to be robust to changes in the illumination and it should avoid detecting non-stationary background objects such as moving leaves, rain, snow, and shadows. In addition, the internal background model should quickly respond to changes in background such as objects that start to move or stop.

We present a new algorithm for video segmentation that processes the input video sequence as a 3D matrix where the third axis is the time domain. Our approach identifies the background by reducing the input dimension using the *diffusion bases* methodology. Furthermore, we describe an iterative method for extracting and deleting the background. The algorithm has two versions and thus covers the complete range of backgrounds: one for scenes with static backgrounds and the other for scenes with dynamic (moving) backgrounds.

keywords: Video Segmentation, Background subtraction, Markov processes, Graph algorithms.

*Corresponding author: alonschc@mta.ac.il, Tel:+972-54-5456226, Fax:+972-3-6803342

1 Introduction

Video surveillance systems, tracking systems, statistical packages that count people, games, etc. seek to automatically identify people, objects, or events of interest in different environment types. Typically, these systems consist of stationary cameras, that are directed at offices, parking lots, playgrounds, fences and so on, together with computer systems that process the video frames. Human operators or other processing elements are notified about salient events. There are many needs for automated surveillance systems in commercial, law enforcement, and military applications. In addition to the obvious security applications, video surveillance technology has been proposed to measure traffic flow, detect accidents on highways, monitor pedestrian congestion in public spaces, compile consumer demographics in shopping malls and amusement parks, log routine maintenance tasks at nuclear facilities, and count endangered species. The numerous military applications include patrolling national borders, measuring the flow of refugees in troubled areas, monitoring peace treaties, and providing secure perimeters around bases.

Substraction of backgrounds, which are captured by static cameras, can be useful to achieve low-bit rate video compression for transmission of rich multimedia content. The subtracted background is transmitted once, followed by the segmented objects which are detected.

A common element in surveillance systems is a module that performs background subtraction to distinguish between background pixels, which should be ignored, and foreground pixels, which should be processed for identification or tracking. The difficulty in background subtraction is not to differentiate, but to maintain the background model, its representation and its associated statistics. In particular, capturing the background in frames where the background can change over time. These changes can be moving trees, leaves, water flowing, sprinklers, fountains, video screens (billboards) just to name a few typical examples. Other forms of changes are weather changes like rain and snow, illumination changes like turning on and off the light in a room and changes in daylight. We refer to this background type as *dynamic background* (DBG) while a background without changes or with slight changes is referred to as *static background* (SBG).

In this paper, we present a new method for capturing the background. It is based on the application of the *diffusion bases* (DB) algorithm. Moreover, we develop real time iterative method for background subtraction in order to separate between background and foreground pixels while

overcoming the presence of changes in the background. The main steps of the algorithm are:

- Extract the background frame by dimensionality reduction via the application of the DB algorithm.
- Subtract the background from the input sequence.
- Threshold the subtracted sequence.
- Detect the foreground objects by applying *depth first search* (DFS).

We propose two versions of the algorithm - one for static background and the other for dynamic background. To handle dynamic background, a learning process is applied to data that contains only the background objects in order to generate a frame that extracts the DBG. The proposed algorithm outperform current state-of-the-art algorithms.

The rest of this paper is organized as follows: In section 2, related algorithms for background subtraction are presented. In section 3, we present the the *diffusion bases* (DB) algorithm. The main algorithm, that is called the *background substraction algorithm using diffusion bases* (BSDB), is presented in section 4. In section 5, we present experimental results, a performance analysis of the BSDB algorithm and we compare it to other background subtraction algorithms.

2 Related work

Background subtraction is a widely used approach for detection of moving objects in video sequences that are captured by static cameras. This approach detects moving objects by differentiating between the current frame and a reference frame, often called the background frame, or background model. In order to extract the objects of interest, a threshold can be applied on the subtracted frame. The background frame should faithfully represent the scene. It should not contain moving objects. In addition, it must be regularly updated in order to adapt to varying conditions such as illumination and geometry changes. This section provides a review of the current state-of-the-art background subtraction techniques. These techniques range from simple approaches,

aiming to maximize speed and minimizing the memory requirements, to more sophisticated approaches, aiming to achieve the highest possible accuracy under any possible circumstances. The goal of these approaches is to run in real-time. Additional references can be found in [1, 2, 3].

Temporal median filter -

In [4], is was proposed to use the median value of the last n frames as the background model. This provides an adequate background model even if the n frames are subsampled with respect to the original frame rate by a factor of 10 [5]. The median filter is computed on a special set of values that contains the last n subsampled frames and the last computed median value. This combination increases the stability of the background model [5].

A fundamental shortcoming of the the median-based approach is the need to store the recent pixel values in order to facilitate the median computation. Moreover, the median filter can not be described by rigorous statistics and does not provide a deviation measure with which the subtraction threshold can be adapted.

Gaussian average -

This approach models the background independently at each pixel location (i, j) [6]. The model is based on ideally fitting a Gaussian probability density function (pdf) to the last n pixels. At each new frame at time t, a running average is computed by $\psi_t = \alpha I_t + (1-\alpha)\psi_{t-1}$ where I_t is the current frame, ψ_{t-1} is the previous average and α is an empirical weight that is often chosen as a tradeoff between stability and quick update.

In addition to speed, the advantage of the running average is given by a low memory requirement. Instead of a buffer with the last n pixel values, each pixel is classified using two parameters (ψ_t, σ_t), where σ_t is the standard deviation. Let $p^t_{i,j}$ be the (i, j) pixel at time t. $p^t_{i,j}$ is classified as a foreground pixel if $|p^t_{i,j} - \psi_{t-1}| > k\sigma_t$. Otherwise $p^t_{i,j}$ is classified as background pixel.

Mixture of Gaussians -

In order to cope with rapid changes in the background, a multi-valued background mode was suggested in [7]. In this model, the probability of observing a certain pixel x at time t is represented by a mixture of k Gaussians distributions: $P(x_t) = \Sigma_{i=1}^{k} w_{i,t}\eta(x_t, \psi_{i,t}, \Sigma_{i,t})$

where for each i-th Gaussian in the mixture at time t, w estimates what portion of the data is accounted for by this Gaussian, ψ is the mean value, Σ is the covariance matrix and η is a Gaussian probability density function. In practice, k is set to be between 3 and 5.

Each of the k Gaussian distributions describes only one of the observable background or foreground objects. The distributions are ranked according to the ratio between their peak amplitude w_i and their standard deviation σ_i. Let Th be the threshold value. The first B distributions that satisfy $\Sigma_{i=1}^{B} w_i > Th$ are accepted as background. All the other distributions are considered as foreground.

Let I_t be a frame at time t. At each frame I_t, two events take place simultaneously: assigning the new observed value x_t to the best matching distribution and estimating the updated model parameters. The distributions are ranked and the first that satisfies $(x_t - \psi_{i,t})/\sigma_{i,t} > 2.5$ is a match for x_t.

Kernel density estimation (KDE) -

This approach models the background distribution by a non-parametric model that is based on a Kernel Density Estimation (KDE) of the buffer of the last *n* background values ([8]). KDE guarantees a smooth, continuous version of the histogram of the most recent values that are classified as background values. This histogram is used to approximate the background pdf.

The background pdf is given as a sum of Gaussian kernels centered at the most recent n background values, x_t: $P(x_t) = \frac{1}{n}\Sigma_{i=1}^{n}\eta(x_t - x_i, \Sigma_t)$ where η is the kernel estimator function and Σ_t represents the kernel function bandwidth. Σ is estimated by computing the median absolute deviation over the sample for consecutive intensity values of the pixel. Each Gaussian describes just one sample data. The buffer of the background values is selectively updated in a FIFO order for each new frame I_t.

In this application two similar models are concurrently used, one for long-term memory and the other for short-term memory. The long-term model is updated using a *blind* update mechanism that prevents incorrect classification of background pixels.

Sequential kernel density approximation -

Mean-shift vector techniques have been proved to be an effective tool for solving a variety of pattern recognition problems e.g. tracking and segmentation ([9, 10]). One of the main advantages of these techniques is their ability to directly detect the main modes of the pdf from the sample data while relying on a minimal set of assumptions. Unfortunately, the computational cost of this approach is very high. As such, it is not immediately applicable to modeling background pdfs at the pixel level.

To solve this problem, computational optimizations are used to mitigate the computational high cost ([11]). Moreover, the mean-shift vector can be used only for an off-line model initialization [12], i.e. the initial set of Gaussian modes of the background pdf is detected from an initial sample set. The real-time model is updated by simple heuristics that handle mode adaptation, creations, and merging.

Co-occurrence of image variations -

This method exploits spatial cooccurrences of image variations ([13]). It assumes that neighboring blocks of pixels that belong to the background should have similar variations over time. The disadvantage of this method is that it does not handle blocks at the borders of distinct background objects.

This method divides each frame to distinct blocks of $N \times N$ pixels where each block is regarded as an N^2-component vector. This trades-off resolution with high speed and better stability. During the learning phase, a certain number of samples is acquired at a set of points, for each block. The temporal average is computed and the differences between the samples and the average, called the *image variations*, is calculated. Then the $N^2 \times N^2$ covariance matrix is computed with respect to the average. An eigenvector transformation is applied to reduce the dimensions of the image variations.

For each block b, a classification phase is performed: the corresponding current eigen-*image-variations* are computed on a neighboring block of b. Then the image variation is expressed as a linear interpolation of its L-nearest neighbors in the eigenspace. The same interpolation coefficients are applied on the values of b, to provide an estimate for its current eigen-*image-variations*.

Eigen-backgrounds -

This approach is based on an eigen-decomposition of the whole image [14]. During a learning phase, samples of n images are acquired. The average image is then computed and subtracted from all the images. The covariance matrix is computed and the best eigenvectors are stored in an eigenvector matrix. For each frame I, a classification phase is executed: I is projected onto the eigenspace and then projected back onto the image space. The output is the background frame, which does not contain any small moving objects. A threshold is applied on the difference between I and the background frame.

3 Dimensionality reduction

Dimensionality reduction has been extensively researched. Classic techniques for dimensionality reduction such as Principal Component Analysis (PCA) and Multidimensional Scaling (MDS) are simple to implement and can be efficiently computed. However, they guarantee to discover the true structure of a data set only when the data set lies on or near a linear subspace of the high-dimensional input space ([15]). These methods are highly sensitive to noise and outliers since they take into account the distances between *all* pairs of points. Furthermore, PCA and MDS fail to detect non-linear structures.

More recent dimensionality reduction methods like Local Linear Embedding (LLE) [16] and ISOMAP [17] amend this pitfall by considering for each point only the distances to its closest neighboring points in the data. Recently, Coifman and Lafon [18] introduced the *Diffusion Maps* (DM) algorithm which is a manifold learning scheme. DM embeds high dimensional data into an Euclidean space of substantially smaller dimension while preserving the geometry of the data set. The global geometry is preserved by maintaining the local neighborhood geometry of each point in the data set. DM uses a random walk distance that is more robust to noise since it averages all the paths between a pair of points.

Diffusion Bases (DB) - a dual algorithm to the DM algorithm - is described in (A. Schclar and A. Averbuch. "Segmentation and anomalies detection in hyper-spectral images via diffusion bases", preprint, 2008). The DB algorithm is dual to the DM algorithm in the sense that it explores the variability among the *coordinates* of the original data. Both algorithms share a graph Laplacian

construction, however, the DB algorithm uses the Laplacian eigenvectors as an orthonormal system on which it projects the original data.

3.1 Diffusion Bases (DB)

This section reviews the DB algorithm for dimensionality reduction. Let $\Omega = \{x_i\}_{i=1}^{m}$, $x_i \in \mathbb{R}^n$, be a data set and let $x_i(j)$ denote the j^{th} coordinate of x_i, $1 \leq j \leq n$. We define the vector $y_j \triangleq (x_1(j), \ldots, x_m(j))$ as the vector whose components are composed of the j^{th} coordinate of all the points in Ω. The DB algorithm consists of the following steps:

- Construct the data set $\Omega' = \{y_j\}_{j=1}^{n}$
- Build a non-directed graph G whose vertices correspond to Ω' with a non-negative and fast-decaying weight function w_ε that corresponds to the *local* point-wise similarity between the points in Ω'. By fast decay we mean that given a scale parameter $\varepsilon > 0$ we have $w_\varepsilon(y_i, y_j) \to 0$ when $\|y_i - y_j\| \gg \varepsilon$ and $w_\varepsilon(y_i, y_j) \to 1$ when $\|y_i - y_j\| \ll \varepsilon$. One of the common choices for w_ε is
$$w_\varepsilon(y_i, y_j) = \exp\left(-\frac{\|y_i - y_j\|^2}{\varepsilon}\right) \tag{1}$$
where ε defines a notion of neighborhood by defining a ε-neighborhood for every point y_i.
- Construction of a random walk on the graph G via a Markov transition matrix P. P is the row-stochastic version of w_ε which is derived by dividing each row of w_ε by its sum[1].
- Perform an eigen-decomposition of P to produce the left and the right eigenvectors of P: $\{\psi_k\}_{k=1,\ldots,n}$ and $\{\xi_k\}_{k=1,\ldots,n}$, respectively. Let $\{\lambda_k\}_{k=1,\ldots,n}$ be the eigenvalues of P where $|\lambda_1| \geq |\lambda_2| \geq \ldots \geq |\lambda_n|$.
- The right eigenvectors of P constitute an orthonormal basis $\{\xi_k\}_{k=1,\ldots,n}$, $\xi_k \in \mathbb{R}^n$. These eigenvectors capture the *non-linear* coordinate-wise variability of the original data.
- Next, we use the spectral decay property of the spectral decomposition to extract only the first η eigenvectors $BS \triangleq \{\xi_k\}_{k=1,\ldots,\eta}$, which contain the *non-linear* directions with the highest variability of the coordinates of the original data set Ω.

[1] P and the graph Laplacian $I - P$ (see [19]) share the same eigenvectors.

- We project the original data Ω onto the basis BS. Let Ω_{BS} be the set of these projections: $\Omega_{BS} = \{g_i\}_{i=1}^{m}$, $g_i \in \mathbb{R}^{\eta}$, where $g_i = (x_i \cdot \xi_1, \ldots, x_i \cdot \xi_\eta)$, $i = 1, \ldots, m$ and $\cdot$ denotes the inner product operator. Ω_{BS} contains the coordinates of the original points in the orthonormal system whose axes are given by BS. Alternatively, Ω_{BS} can be interpreted in the following way: the coordinates of g_i contain the correlation between x_i and the directions given by the vectors in BS.

A summary of the DB procedure is given in Algorithm 1. An enhancement of the spectral decomposition is described in (A. Schclar and A. Averbuch. "Segmentation and anomalies detection in hyper-spectral images via diffusion bases", preprint, 2008).

4 The Background Subtraction Algorithm using Diffusion Bases (BSDB)

In this section we present the BSDB algorithm. The algorithm has two versions:

Static background subtraction using DB (SBSDB): We assume that the background is static (SBG) – see section 4.1. The video sequence is captured on-line. Gray level images are sufficient for the processing.

Dynamic background substraction using DB (DBSDB): We assume that the background is moving (DBG) – see section 4.2. This algorithm uses off-line (training) and on-line (detection) procedures. As opposed to the SBSDB, this algorithm requires color (RGB) frames.

We assume that in both algorithms the camera is static.

4.1 Static background subtraction algorithm using DB (SBSDB)

In this section we describe the on-line algorithm that is applied on a video sequence that is captured by a static camera. We assume that the background is static. The SBSDB algorithm captures the static background, subtracts it from the video sequence and segments the subtracted output.

Algorithm 1 The Diffusion Basis algorithm.

DiffusionBasis(Ω', w_ε, ε, η)

1. Calculate the weight function $w_\varepsilon\left(y_i, y_j\right),\ i, j = 1, \ldots n$, (Eq. 1).

2. Construct a Markov transition matrix P by normalizing the sum of each row in w_ε to be 1:

$$p\left(y_i, y_j\right) = \frac{w_\varepsilon\left(y_i, y_j\right)}{d\left(y_i\right)}$$

 where $d\left(y_i\right) = \sum_{j=1}^{n} w_\varepsilon\left(y_i, y_j\right)$.

3. Perform a spectral decomposition of P

$$p\left(y_i, y_j\right) \equiv \sum_{k=1}^{n} \lambda_k \xi_k\left(y_i\right) \psi_k\left(y_j\right)$$

 where the left and the right eigenvectors of P are given by $\{\psi_k\}$ and $\{\xi_k\}$, respectively, and $\{\lambda_k\}$ are the eigenvalues of P in descending order of magnitude.

4. Let $BS \triangleq \{\xi_k\}_{k=1,\ldots,\eta}$.

5. Project the original data Ω onto the orthonormal system BS:

$$\Omega_{BS} = \{g_i\}_{i=1}^{m},\ g_i \in \mathbb{R}^\eta$$

 where

$$g_i = \left(x_i \cdot \xi_1, \ldots, x_i \cdot \xi_\eta\right),\ i = 1, \ldots, m,\ \xi_k \in BS,\ 1 \leq k \leq \eta$$

 and $\cdot$ is the inner product.

6. **return** Ω_{BS}.

The input to the algorithm is a sequence of video frames in gray-level format. The algorithm produces a binary mask for each video frame. The pixels in the binary mask that belong to the background are assigned 0 values while the other pixels are assigned to be 1.

4.1.1 Off-line algorithm for capturing static background

In order to capture the static background of a scene, we reduce the dimensionality of the input sequence by applying the DB algorithm (Algorithm 1 in section 3.1). The input to the algorithm consists of n frames that form a datacube.

Formally, let $D_n = \left\{ s_{i,j}^t, i,j = 1,...,N, t = 1,...,n \right\}$ be the input datacube of n frames each of size $N \times N$ where $s_{i,j}^t$ is the pixel at position (i,j) in the video frame at time t. We define the vector $P_{i,j} \triangleq \left(s_{i,j}^1, ..., s_{i,j}^n \right)$ to be the values of the $(i,j)^{th}$ coordinate at all the n frames in D_n. This vector is referred to as a *hyperpixel*. Let $\Omega_n = \{P_{i,j}\}$, $i,j = 1,...,N$ be the set of all hyperpixels. We define $F_t \triangleq (s_{1,1}^t, ..., s_{N,N}^t)$ to be a 1-D vector representing the video frame at time t. We refer to F_t as a frame-vector. Let $\Omega_n' \triangleq \{F_t\}_{t=1}^n$ be the set of all frame-vectors.

We apply the DB algorithm to Ω_n in Ω_{BS}=DiffusionBasis(Ω_n', w_ε, ε, η) where w_ε is defined by Eq. 1, ε and η are defined in section 3.1 - see Algorithm 1. The output is the projection of every hyperpixel on the diffusion basis which embeds the original data D_n into a reduced space. The first vector of Ω_{BS} represents the background of the input frames. Let $bg_V = (x_i)$, $i = 1,...,N^2$ be this vector. We reshape bg_V into the matrix $bg_M = (x_{i,j})\,, i,j = 1,...,N$. Then, bg_M is normalized to be between 0 to 255. The normalized background is denoted by $\widehat{bg}_M$.

4.1.2 On-line algorithm for capturing a static background

In order to make the algorithm suitable for on-line applications, the incoming video sequence is processed by using a *sliding window* (SW) of size m. Thus, the number of frames that are input to the algorithm is m. Naturally, we seek to minimize m in order to obtain a faster result from the algorithm. We found empirically that the algorithm produces good results for values of m as low as $m = 5, 6$ and 7. The delay of 5 to 7 frames is negligible and renders the algorithm to be suitable for on-line applications.

Let $S = (s_1, ..., s_i..., s_m, s_{m+1}, ..., s_n)$ be the input video sequence. we apply the algorithm that is described in section 4.1.1 to every SW. The output is a sequence of background frames

$$\widehat{BG} = \left((\widehat{bg}_M)_1, ..., (\widehat{bg}_M)_i, ..., (\widehat{bg}_M)_m, (\widehat{bg}_M)_{m+1}, ..., (\widehat{bg}_M)_n \right) \tag{2}$$

where $(\widehat{bg}_M)_i$ is the background that corresponds to frame s_i and $(\widehat{bg}_M)_{n-m+2}$ till $(\widehat{bg}_M)_n$ are equal to $(\widehat{bg}_M)_{n-m+1}$. Figure 1 describes how the SW is shifted.

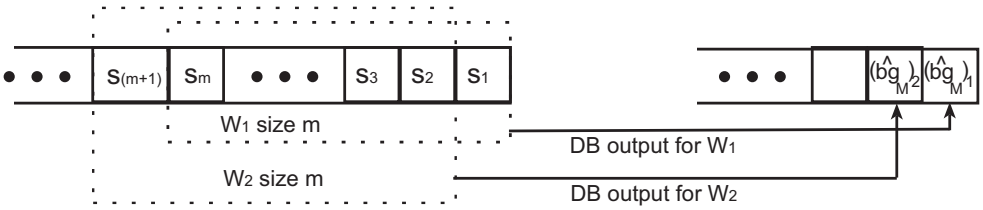


Figure 1: Illustration of how the SW is shifted. $W_1 = (s_1, ..., s_m)$ is the SW for s_1. $W_2 = (s_2, ..., s_{m+1})$ is the SW for s_2, etc. The backgrounds of s_i and s_{i+1} are denoted by $(\widehat{bg}_M)_i$ and $(\widehat{bg}_M)_{i+1}$, $i = 1, ..., n - m + 1$, respectively.

The SW results in a faster execution time of the DB algorithm. The weight function w_ε (Eq. 1) is not recalculated for all the frame in the SW. Instead, w_ε is only updated according to the new frame that enters the SW and the one that exits the SW. Specifically, let $W_t = (s_t, ..., s_{t+m-1})$ be the SW at time t and let $W_{t+1} = (s_{t+1}, ..., s_{t+m})$ be the SW at time $t + 1$. At time $t + 1$, w_ε is calculated only for s_{t+m} and the entries that correspont to s_t are removed from w_ε.

4.1.3 The SBSDB algorithm

The SBSDB on-line algorithm captures the background of each SW according to section 4.1.2. Then it subtracts the background from the input sequence and thresholds the output to get the background binary mask.

Let $S = (s_1..., s_n)$ be the input sequence. For each frame $s_i \in S$, $i = 1, ..., n$, we do the following:

- Let $W_i = (s_i, ..., s_{i+m-1})$ be the SW of s_i. The on-line algorithm for capturing the background (section 4.1.2) is applied to W_i. The output is the background frame $(\widehat{bg}_M)_i$.
- The SBSDB algorithm subtracts $(\widehat{bg}_M)_i$ from the original input frame by $\bar{s}_i = s_i - (\widehat{bg}_M)_i$. Then, each pixel in $\bar{s}_i$ that has a negative value is set to 0.
- A threshold is applied to $\bar{s}_i$. The threshold is computed in section 4.1.4. For $k, l = 1, ..., N$ the output is defined as follows:

$$\tilde{s}_i(k,l) = \begin{cases} 0, & \text{if it is a background pixel;} \\ 1, & \text{otherwise.} \end{cases}$$

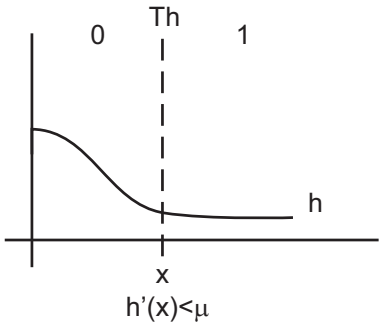


Figure 2: An example how to use the histogram h for finding the threshold value. Th is set to x since $h'(x) < \mu$.

4.1.4 Threshold computation for a grayscale input

The threshold Th, which separates between background and foreground pixels, is calculated in the last step of the SBSDB algorithm. The SBSDB algorithm subtracts the background from the input frame and sets pixels with negative values to zero. Usually, the histogram of a frame after subtraction will be high at small values and low at high values. The SBSDB algorithm smooths the histogram in order to compute the threshold value accurately.

Let h be the histogram of a frame and let μ be a given parameter which provides a threshold for the slope of h. μ is chosen to be the magnitude of the slope where h becomes moderate. We scan h from its global maximum to the right. We set the threshold Th to be the smallest value of x that satisfies $h'(x) < \mu$ where h'_x is the first derivative of h at point x, i.e. the slope of h at point x. The background/foreground classification of the pixels in the input frame $\bar{s}_i$ is determined according to Th. Specifically, for $k, l = 1, ..., N$

$$\tilde{s}_i(k,l) = \begin{cases} 0, & \text{if } \bar{s}_i(k,l) < Th; \\ 1, & \text{otherwise.} \end{cases}$$

Fig. 2 illustrates how to find the threshold.

4.2 Dynamic background subtraction algorithm using DB (DBSDB)

In this section, we describe an on-line algorithm that handles video sequences that are captured by a static camera. We assume that the background is dynamic (moving). The DBSDB applies an

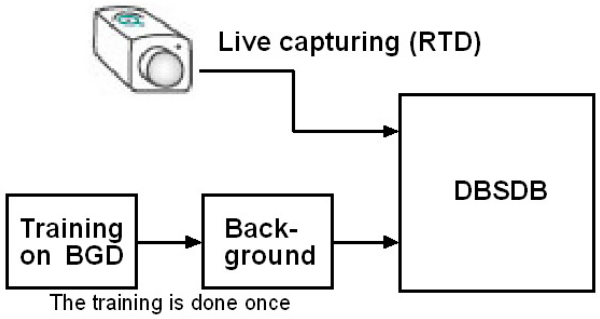


Figure 3: The inputs to the DBSDB algorithm. The training is done once on the BGD. It produces the background which is input to the DBSDB. The RTD is the on-line input to the DBSDB.

off-line procedure that captures the dynamic background and an on-line background subtraction algorithm. In addition, the DBSDB algorithm segments the video sequence after the background subtraction is completed.

The input to the algorithm consists of two components:

Background training data: A video sequence of the scene without foreground objects. This training data can be obtained from the frames in the beginning of the video sequence. This sequence is referred to as the *background data* (BGD).

Data for classification: A video sequence that contains background and foreground objects. The classification of the objects is performed on-line. We refer to this sequence as the *real-time data* (RTD).

For both input components, the video frames are assumed to be in RGB – see Fig. 3.

The algorithm is applied to every video frame and a binary mask is constructed in which the pixels that belong to the background are set to 0 while the foreground pixels are set to 1.

4.2.1 Iterative method for capturing a dynamic background - training

The algorithm that is described in section 4.1.3, does not handle well on-going changes in the background, such as illumination differences between frames, moving leaves, water flowing, etc. In the following, we present a method that is not affected by background changes. An iterative

procedure is applied on the BGD in order to capture the movements in the scene. This procedure constitutes the training step of the algorithm.

Let $B = (b_1, ..., b_m)$ be the BGD input sequence and let bg_M^{final} be the output background frame. bg_M^{final} is initialized to zeros. Each iteration contains the following steps:

- Application of the off-line algorithm (section 4.1.1) in order to capture the static background of B. The BGD is treated as a single sliding window of length m. The output consists of the background frames bg_M and $\widehat{bg}_M$ where $\widehat{bg}_M$ is the normalization of bg_M.
- $\widehat{bg}_M$ is subtracted from each frame in B by $\bar{b}_j = b_j - \widehat{bg}_M$, $j = 1, ..., m$. In case the input is in grayscale format, we set to zero each pixel in $\bar{b}_j$ that has a negative value. The output is the sequence $\bar{B} = \left(\bar{b}_1, ..., \bar{b}_m\right)$.
- bg_M is added to bg_M^{final} by $bg_M^{final} = bg_M^{final} + bg_M$.
- $\bar{B}$ is the input for the next iteration, $B = \bar{B}$.

The iterative process stops when a given number of pixels in B are equal to or smaller than zero. Finally, bg_M^{final} is normalized to be between 0 to 255. The normalized background is denoted by $\widehat{bg}_M^{final}$. The output of this process is composed of bg_M^{final} and $\widehat{bg}_M^{final}$.

4.2.2 The DBSDB algorithm

In this section, we describe the DBSDB algorithm which handles video sequence that contain a dynamic background. The DBSDB algorithm consists of a training phase, which captures the BGD (section 4.2.1), and a classification phase, which is applied on the RTD. Both phases process grayscale and RGB versions of the input and generate grayscale and RGB outputs. The final phase combines the output from the grayscale classification phase and the output from the RGB classification phase.

Formally, let $S^{rgb} = \left(s_1^{rgb}, .., s_n^{rgb}\right)$ and $B^{rgb} = \left(b_1^{rgb}, ..., b_m^{rgb}\right)$ be the RTD, which is the on-line captured video sequence, and the BGD, which is the off-line video sequence for the training phase (section 4.2.1), respectively.

The DBSDB algorithm consists of the following:

1. **The grayscale training phase**

 - Convert B^{rgb} into grayscale format. The grayscale sequence is denoted by B^g.
 - Apply the SBSDB algorithm to B^g *excluding* the threshold computation, as was done in section 4.1.3. The output is a sequence of background frames $\bar{B}^g$.
 - Capture the *dynamic background* (DBG) in $\bar{B}^g$ (section 4.2.1). The output is the background frame given by $\widehat{(bg_M^{final})}^g$.

2. **The RGB training phase:**

 Capture the DBG in each of the RGB channels of B^{rgb} (section 4.2.1). The output is the background frame denoted by $\widehat{(bg_M^{final})}^{rgb}$.

3. **The grayscale classification phase:**

 S^{rgb} is converted into grayscale format. The grayscale sequence is denoted by S^g. The SBSDB algorithm is applied on S^g *excluding* the threshold computation, as it is described in section 4.1.3. This process is performed once or, in some cases, iteratively twice. The output is denoted by $\bar{S}^g$.

 For each frame $\bar{s}_i^g \in \bar{S}^g$, $i = 1, ..., n$, we do the following:

 - $\widehat{(bg_M^{final})}^g$ is subtracted from $\bar{s}_i^g$ by $\tilde{s}_i^g = \bar{s}_i^g - \widehat{(bg_M^{final})}^g$. Then, each pixel in $\tilde{s}_i^g$ that has a negative value is set to 0.
 - A threshold is applied to $\tilde{s}_i^g$. The threshold is computed as in section 4.1.4. The output is set to:

 $$\breve{s}_i^g(k,l) = \begin{cases} 0, & \text{if it is a background pixel;} \\ 1, & \text{otherwise.} \end{cases}$$

 for $k, l = 1, ..., N$.

4. **The RGB classification phase:**

 For each frame $s_i^{rgb} \in S^{rgb}$, $i = 1, ..., n$, we do the following:

 - $\widehat{(bg_M^{final})}^{rgb}$ is subtracted from s_i^{rgb} by $\bar{s}_i^{rgb} = s_i^{rgb} - \widehat{(bg_M^{final})}^{rgb}$.
 - $\bar{s}_i^{rgb}$ is normalized to be between 0 to 255. The normalized frame is denoted by $\tilde{s}_i^{rgb}$.

- A threshold is applied to $\tilde{s}_i^{rgb}$. The threshold is computed according to section 4.2.2. The output is set to:

$$\breve{s}_i^{rgb}(k,l) = \begin{cases} 0, & \text{if it is a background pixel;} \\ 1, & \text{otherwise.} \end{cases}$$

for $k, l = 1, ..., N$.

5. **The DFS phase:**

This phase combines the $\breve{s}_i^g$ and $\breve{s}_i^{rgb}$ from the grayscale and RGB classification phases, respectively. Since $\breve{s}_i^g$ contains false negative detections (not all the foreground objects are found) and $\breve{s}_i^{rgb}$ contains false positive detections (background pixels are classified as foreground pixels), we use each foreground pixel in $\breve{s}_i^g$ as a reference point from which we begin the application of a DFS on $\breve{s}_i^{rgb}$ (see section 4.2.2).

Threshold computation for RGB input In the last step of the RGB classification phase in the DBSDB algorithm, the thresholds that separate between background pixels and foreground pixels are computed for each of the RGB components. The DBSDB algorithm subtracts the background from the input frame, therefore, the histogram of a frame after the subtraction is high in the center and low at the right and left ends, where the center area corresponds to the background pixels. The DBSDB algorithm smooths the histogram in order to compute the threshold values accurately.

Let h be the histogram and let μ be a given parameter which provides a threshold for the slope of h. μ should be chosen to be the value of the slope where h becomes moderate. We denote the thresholds to be Th^r and Th^l. We scan h from its global maximum to the right. $Th^r = x$ if x is the first coordinate that satisfies $h'(x) < \mu$ where $h'(x)$ denotes the first derivative of h at point x, i.e. the slope of h at point x. We also scan h from its global maximum to the left. $Th^l = y$ if y is the first coordinate that satisfies $h'(y) > -\mu$.

The classification of the pixels in the input frame $\tilde{s}_i^{rgb}$ is determined according to Th^r and Th^l. For each color component and for each $k, l = 1, ..., N$

$$\breve{s}_i^{rgb}(k,l) = \begin{cases} 0, & \text{if } Th^l < \tilde{s}_i^{rgb}(k,l) < Th^r; \\ 1, & \text{otherwise.} \end{cases}$$

See Fig.4 for an example how the thresholds are computed.

The process is executed three times, one for each of the RGB channels. The outputs are combined by a pixel-wise OR operation.

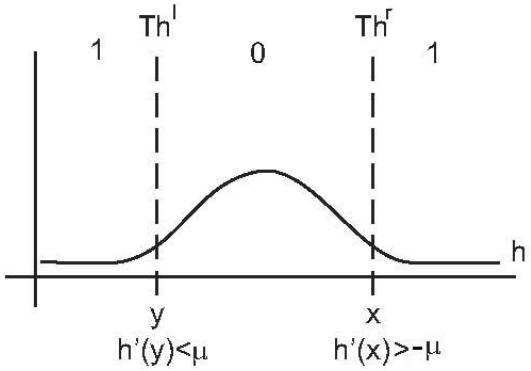


Figure 4: An example that uses the histogram h for finding the threshold values.

Scan by depth-first search (DFS) The last phase of the DBSDB algorithm is the application of a DFS. Let $\breve{s}_i^g \in \breve{S}^g$ and $\breve{s}_i^{rgb} \in \breve{S}^{rgb}$ be the i^{th} output frames of the grayscale and the RGB classification phases, respectively. Each frame is a binary mask represented by a matrix. The DFS phase combines both outputs. In $\breve{s}_i^g$ there are false negative detections and in $\breve{s}_i^{rgb}$ there are false positive detections. We use each foreground pixel in $\breve{s}_i^g$ as a reference point from which we begin a DFS in $\breve{s}_i^{rgb}$. The goal is to find the connected components of the graph whose vertices are constructed from the pixels in $\breve{s}_i^{rgb}$ and whose edges are constructed according to the 8-neighborhood of each pixel.

The graph is constructed as follows:

- A pixel $\breve{s}_i^{rgb}(k,l)$ is a root if $\breve{s}_i^g(k,l)$ is a foreground pixel and it has not been classified yet as a foreground pixel by the algorithm.
- A pixel $\breve{s}_i^{rgb}(k,l)$ is a node if it is a foreground pixel and was not marked yet as a root.
- Let $\breve{s}_i^{rgb}(k,l)$ be a node or a root and let $M_{(k,l)}$ be a 3×3 matrix that represent its 8-neighborhood. A pixel $\breve{s}_i^{rgb}(q,r) \in M_{(k,l)}$ is a child of $\breve{s}_i^{rgb}(k,l)$ if $\breve{s}_i^{rgb}(q,r)$ is a node (see Fig.5).

The DFS is applied from each root in the graph. Each node, that is scanned by the DFS,

represents a pixel that belongs to the foreground objects that we wish to find. The scanned pixels are marked as the new foreground pixels and the others as the new background pixels.

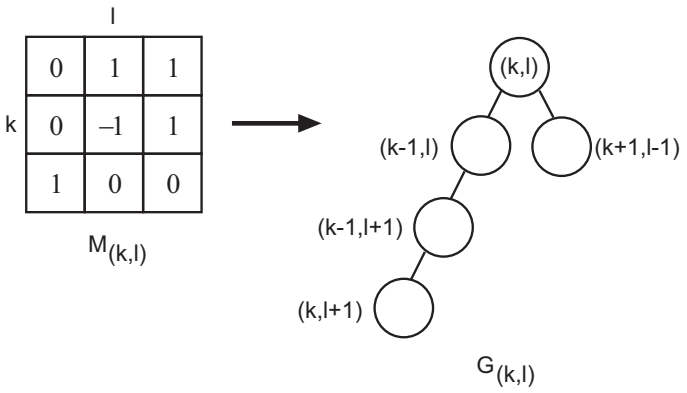


Figure 5: $G_{(k,l)}$ is a graph representation of the neighboring matrix $M_{(k,l)}$ of the root pixel $\breve{s}_i^{rgb}(k,l)$. A root pixel is set to -1, a foreground pixel is set to 1 and a background pixel is set to 0.

4.3 A parallel extension of the SBSDB and the DBSDB algorithms

We propose parallel extensions to the SBSDB and the DBSDB algorithms. We describe this scheme for the SBSDB algorithm and the same scheme can be used for the DBSDB algorithm.

First, the data cube $D_n = \left\{s_{i,j}^t, i,j = 1,...,N, t = 1,...,n\right\}$ is decomposed into overlapping blocks $\{\beta_{k,l}\}$. Next, the SBSDB algorithm is independently applied on each block. This step can run in *parallel*. The final result of the algorithm is constructed using the results from each block. Specifically, the result from each block is placed at its original location in D_n. The result for pixels that lie in overlapping areas between adjacent blocks is obtaind by applying a logical OR operation on the corresponding blocks results.

5 Experimental results

In this section, we present the results from the application of the SBSDB and DBSDB algorithms. The section is divided into three parts: The first part is composed from the results of the SBSDB algorithm when applied to a SBG video. The second part contains the results from the application of the DBSDB algorithm to a DBG video. In the third part we compare between the results obtained



Figure 6: The frames that W_s contains. The test frame s is the top-left frame. The frames are ordered from top-left to bottom-right.

by our algorithm and those obtained by five other background-subtraction algorithms.

5.1 Performance analysis of the SBSDB algorithm

We apply the SBSDB algorithm to a video sequence that consists of 190 grayscale frames of size 256×256. The video sequence was captured by a static camera and is in AVI format with a frame rate of 15 fps. The video sequence shows moving cars over a static background. We apply the sequential version of the algorithm where the size of the SW is set to 5. We also apply the parallel version of the algorithm where the video sequence is divided to four blocks in a 2×2 formation. The overlapping size between two (either horizontally or vertically) adjacent blocks is set to 20 pixels and the size of the SW is set to 10. Let s be the test frame and let W_s be the SW starting at s. In Fig. 6 we show the frames that W_s contains. The output of the SBSDB algorithm for s is shown in Fig. 7.

5.2 Performance analysis of the DBSDB algorithm

We apply the DBSDB algorithm to five video sequences. The first four video sequences are in AVI format with a frame rate of 30 fps. The last video sequence is in AVI format with a frame rate



Figure 7: (a) The background for the test frame s. (b) The test frame s after the subtraction of the background. (c) The output for the test frame s. (d) The output for the test frame s from the parallel version of the algorithm.

of 24 fps. All the video sequences, except the first video sequence, are in RGB format and are of size 320×240. The first video sequence is of size 210×240 and is in RGB format. The video sequences were produced by a static camera and contain a dynamic background.

The input video sequences are:

1. People walking in front of a fountain. It contains moving objects in the background such as water flowing, waving trees and a video screen whose content changes over time. The DBSDB input is a RTD that contains 170 frames and a BGD that contains 100 frames. The output of the DBSDB is presented in Fig. 8(g).

2. A person walking in front of bushes with waving leaves. The DBSDB input is a RTD that contains 88 frames and a BGD that contains 160 frames. The output of the DBSDB algorithm is presented in Fig. 9(g).

3. A moving ball in front of waving trees. The DBSDB input is a RTD that contains 88 frames and a BGD that contains 160 frames. The output of the DBSDB algorithm is presented in Fig. 10. Figure 10(d) contains the result of the sequential version of the algorithm and Fig. 10(g) contains the results of the parallel version. In results of the parallel version the video sequence was divided to four blocks in a 2×2 formation. The overlapping size between two (either horizontally or vertically) adjacent blocks was set to 20 pixels and the size of SW was set to 30.

4. A ball jumping in front of trees and a car passing behind the trees. The DBSDB input is

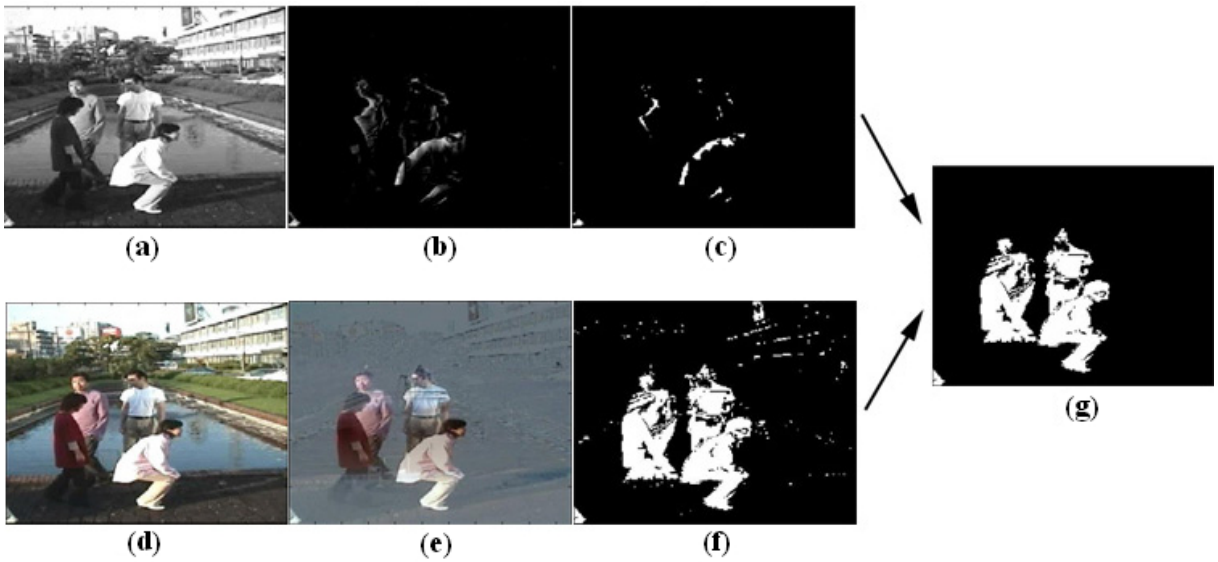


Figure 8: (a), (d) The original test frames in grayscale and RGB, respectively. (b), (e) The grayscale and RGB test frames after the background subtraction in the classification phase of the DBSDB algorithm, respectively. (c), (f) Results after the thresholding of (b) and (e), respectively. (g) The final output of the DBSDB algorithm after the application of the DFS.

a RTD that contains 106 frames and a BGD that contains 160 frames. The output of the DBSDB algorithm is presented in Fig. 10(e).

5. A person walking in front of a sprinkler. The DBSDB input is a RTD that contains 121 frames and a BGD that contains 100 frames. The output of the DBSDB algorithm is presented in Fig. 10(f).

5.3 Performance comparison between the BSDB algorithm and other algorithms

We compared between the BSDB algorithm and five different background subtraction algorithms. The input data and the results are taken from [20]. All the test sequences were captured by a camera that has three CCD arrays. The frames are of size 160x120 in RGB format and are sampled at 4Hz . The test frame that was segmented, in video sequences where the background changes, is taken to be the frame that appears 50 frames after the frame where the background changes. On every

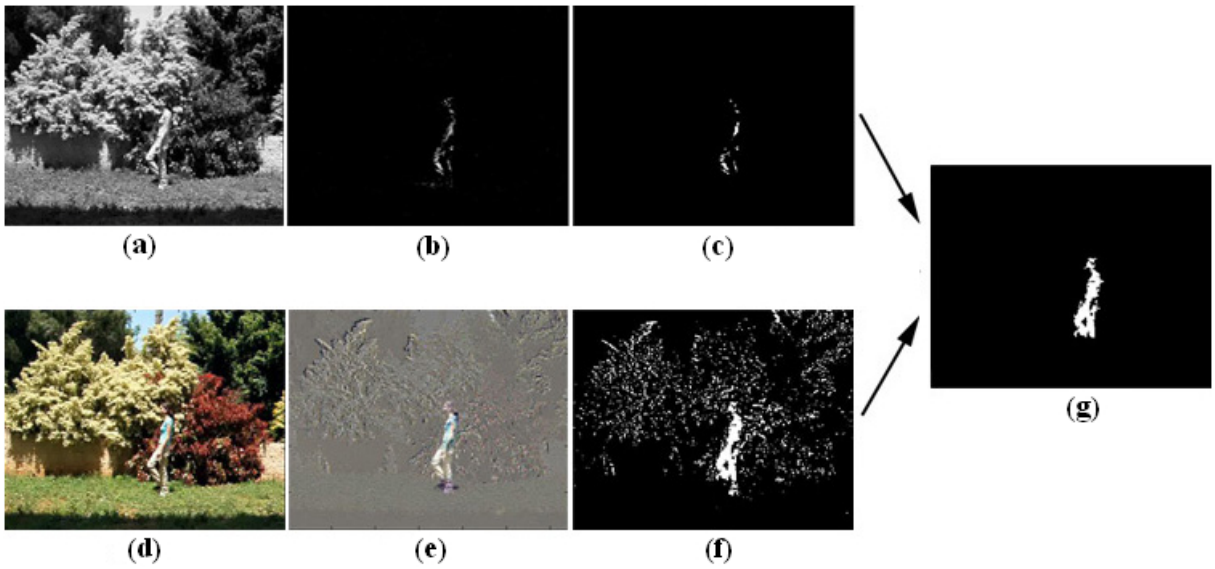


Figure 9: (a), (d) The original test frames in grayscale and RGB, respectively. (b), (e) The grayscale and RGB test frames after the background subtraction in the classification phase of the DBSDB algorithm, respectively. (c), (f) Results after the thresholding of (b) and (e), respectively. (g) The final output of the DBSDB algorithm after the application of the DFS.

output frame (besides the output of the BSDB algorithm), a speckle removal [20] was applied to eliminate islands of 4-connected foreground pixels that contain less than 8 pixels. All other parameters were adjusted for each algorithm in order to obtain visually optimal results over the entire dataset. The parameters were used for all sequences. Each test sequence begins with at least 200 background frames that were used for training the algorithms, except for the bootstrap sequence. Objects such as cars, which might be considered foreground in some applications, were deliberately excluded from the sequences.

Each of the sequences poses a different problem in background maintenance. The chosen sequences and their corresponding problems are:

Background object is moved - Problem: A background object can move. These objects should not be considered as part of the foreground. The sequence contains a person that walks into a conference room, makes a telephone call, and leaves with the phone and a chair in a different position. The test frame is the one that appears 50 frames after the person has left the scene.

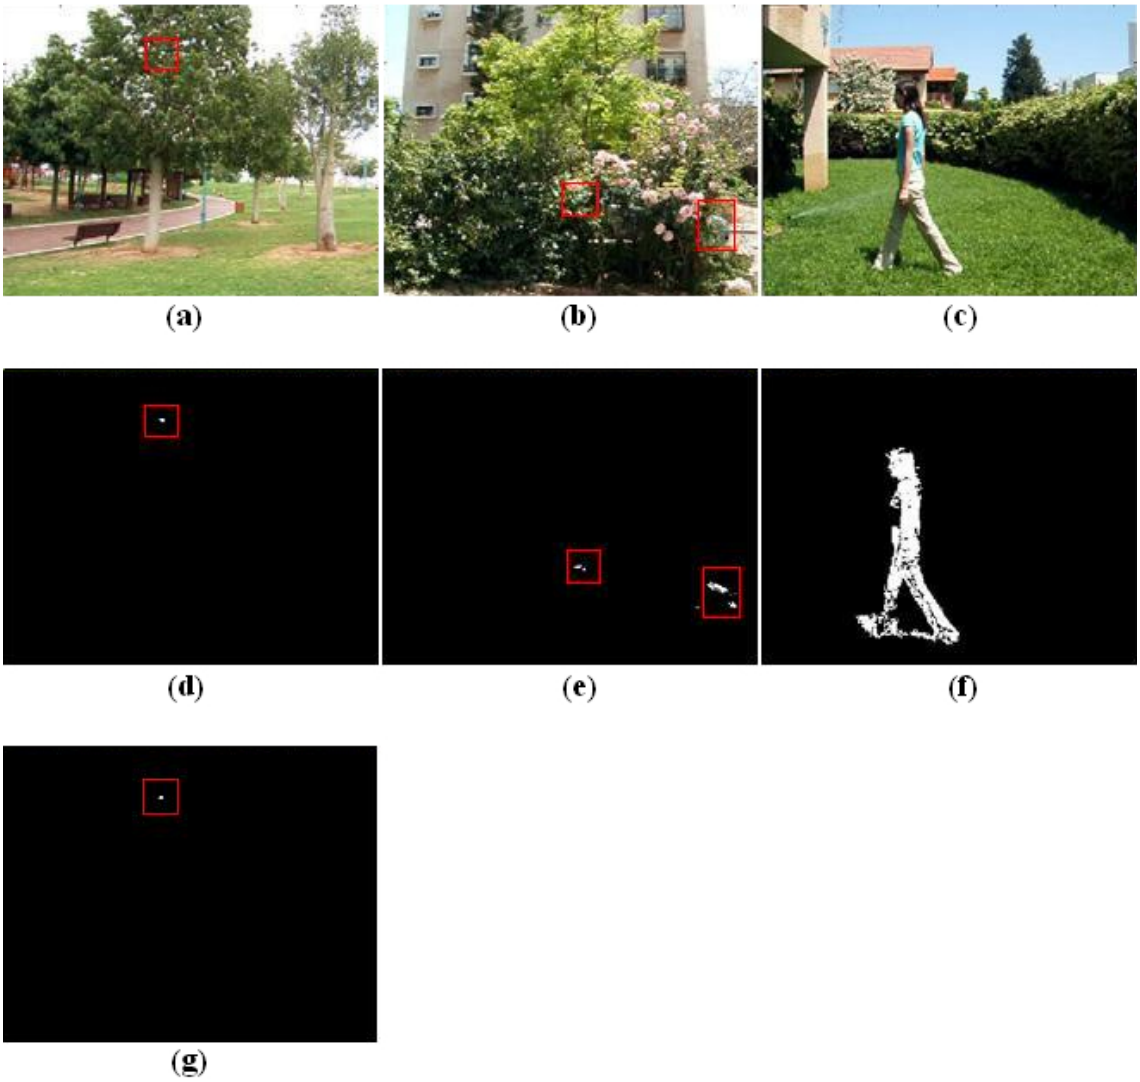


Figure 10: (a)-(c) The original test frames. (d)-(g) The segmented outputs from the application of the DBSDB algorithm. (a) A ball in front of waving trees. (d), (g) The result of the sequential and parallel versions of the algorithm applied on (a), respectively. (b), (e) A ball jumping in front of a tree and a car passing behind the trees. (c), (f) A person walking in front of a sprinkler.

Bootstrapping - Problem: A training period without foreground objects is not available. The sequence contains an overhead view of a cafeteria. There is constant motion and every frame contains people.

Waving Trees - Problem: Backgrounds can contain moving objects. The sequence contains a person walking in front of a swaying tree.

Camouflage - Problem: Pixels of foreground objects may be falsely recognized as background pixels. The sequence contains a monitor on a desk with rolling interference bars. A person walks into the scene and stands in front of the monitor.

We apply six background subtraction algorithms on these sequences, including the algorithm that is presented in this paper. The background subtraction algorithms are:

Adjacent Frame Difference - Each frame is subtracted from the previous frame in the sequence. Absolute differences greater than a threshold are marked as foreground.

Mean and Threshold - Pixel-wise mean values are computed during a training phase, and pixels within a fixed threshold of the mean are considered background.

Mean and Covariance - The mean and covariance of pixel values are updated continuously [21]. Foreground pixels are determined by applying a threshold to the Mahalanobis distance.

Mixture of Gaussians - This algorithm is reviewed in section 2.

Eigen-background - This algorithm is reviewed in section 2.

BSDB - The algorithm presented in this paper (section 4).

The outputs of these algorithms are shown in Fig. 11.

We applied the SBSDB algorithm on the first two video sequences: the moved chair and the bootstrapping. In both cases, the background in the video sequence is static. In the first video sequence, the SBSDB algorithm handles the changes in the position of the chair that is a part of the background. The SBSDB algorithm does not require a training process so it can handle the second video sequence where there is no clear background for training. Algorithms that require a training process can not handle this case.

We applied the DBSDB algorithm on the waving trees and the camouflage video sequences. In both cases, the background in the video sequences is dynamic. In the first video sequence, the DBSDB algorithm captures the movement of the waving trees and eliminates it from the video sequence. The other algorithms produce false positive detections. The DBSDB algorithm does not handle well the last video sequence where the foreground object covers the background moving object (the monitor). In this case the number of false negative detections is significant.

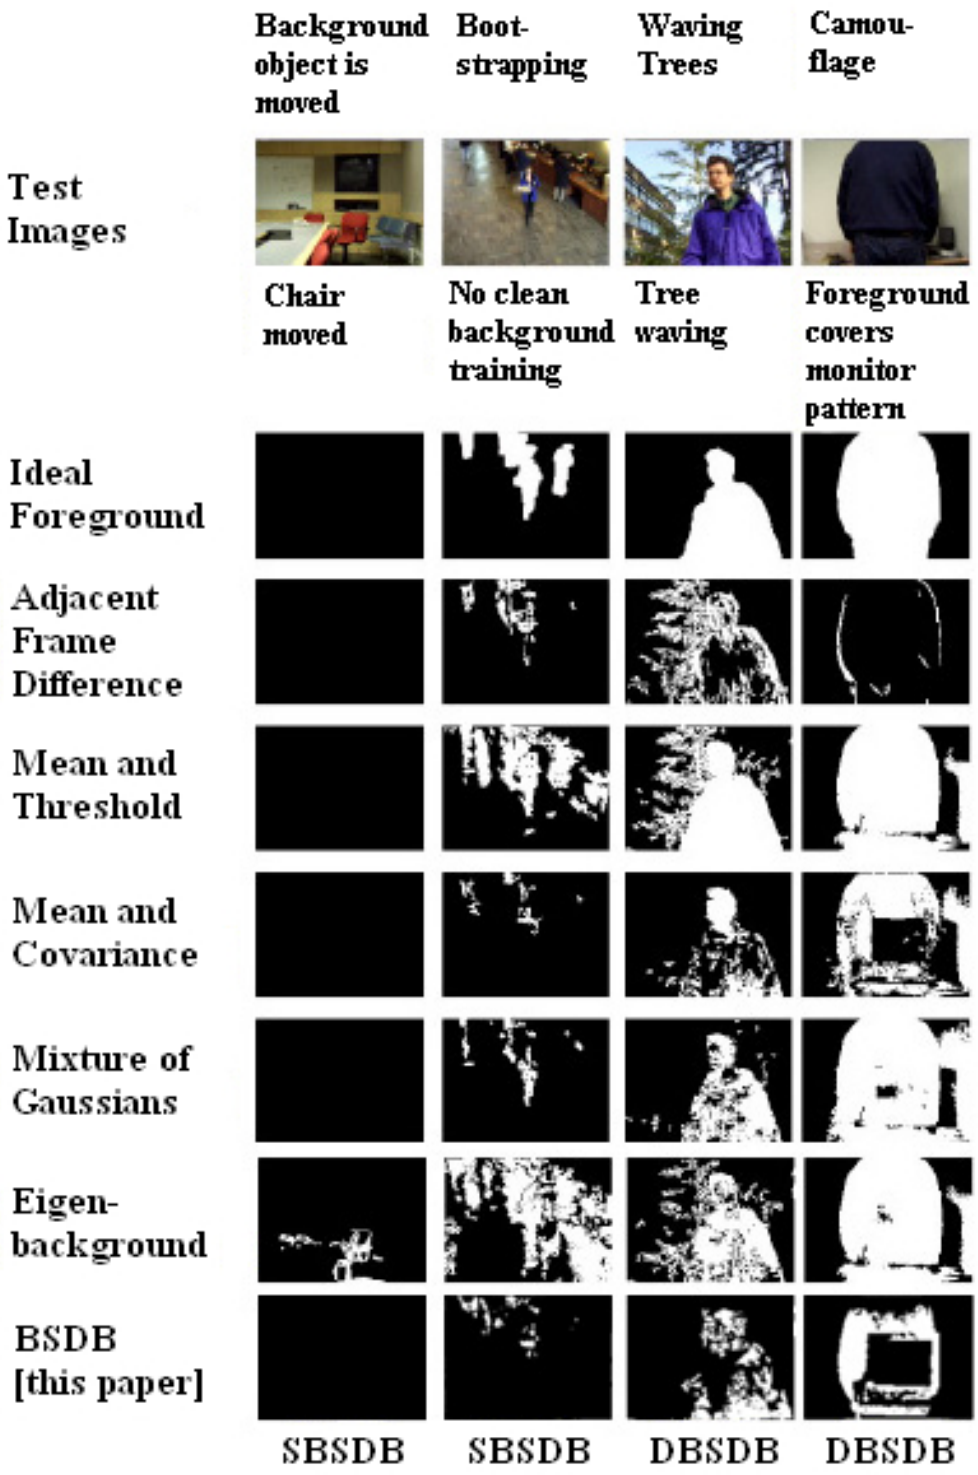


Figure 11: The outputs from the applications of the BSDB and five other algorithms. Each row shows the results of one algorithm, and each column represents one problem in background maintenance. The top row shows the test frames. The second row shows the optimal background outputs.

6 Conclusion and Future Work

We introduced in this work the BSDB algorithm for automatic segmentation of video sequences. The algorithm contains two versions: the SBSDB algorithm for video sequences with static background and the DBSDB algorithm for video sequences that contain dynamic background. The BSDB algorithm captures the background by reducing the dimensionality of the input via the DB algorithm. The SBSDB algorithm uses an on-line procedure while the DBSDB algorithm uses an off-line (training) procedure and an on-line procedure. During the training phase, the DBSDB algorithm captures the dynamic background by iteratively applying the DB algorithm on the background training data. The BSDB algorithm presents a high quality segmentation of the input video sequences. Moreover, it was shown that the BSDB algorithm outperforms the current state-of-the-art algorithms by coping with difficult situations of background maintenance.

The performance of the BSDB algorithm can be enhanced by improving the accuracy of the threshold values. Furthermore, it is necessary to develop a method for automatic computation of μ, which is used in the threshold computation (sections 4.1.4 and 4.2.2).

Additionally, the output of the BSDB algorithm contains a fair amount of false negative detections when a foreground object obscures a brighter background object. This will be improved in future versions of the algorithm.

The BSDB algorithm can be useful to achieve low-bit rate video compression for transmission of rich multimedia content. The captured background is transmitted once followed by the detected segmented objects.

References

[1] R. Collins, A. Lipton, and T. Kanade, "Introduction to the special section on video surveillance," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 745–746, Aug. 2000.

[2] M. Piccardi, "Background subtraction techniques: a review," *IEEE International Confrence on Systems, Man and Cybernetics*, 2004. [Online]. Available: http://www-staff.it.uts.edu.au/~{}massimo/BackgroundSubtractionReview-Piccardi.pdf

[3] A. M. Macivor, "Background subtraction techniques," *Proc. of Image and Vision Computing*, 2000.

[4] B. P. L. Lo and S. A. Velastin, "Automatic congestion detection system for underground platforms," *Proc. ISIMP*, pp. 158–161, May 2001.

[5] R. Cucchiara, C. Grana, M. Piccardi, and A. Prati, "Detecting moving objects, ghosts, and shadows in video streams," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 10, pp. 1337–1442, 2003.

[6] C. Wren, A. Azarhayejani, T. Darrell, and A. P. Pentland, "Pfinder: real-time tracking of the human body," *IEEE Transactions on Pattern Analysis and Machine Analysis*, vol. 19, no. 7, pp. 780–785, 1997.

[7] C. Stauffer and W. E. L. Grimson, "Adaptive background mixture models for real-time tracking," *IEEE Computer Vision and Pattern Recognition*, pp. 246–252, June 1999.

[8] A. Elgammal, D. Hanvood, and L. S. Davis, "Nonparametric model for background subtraction," *Proceedings of the European Conference on Computer Vision*, pp. 751–767, June 2000.

[9] D. Comaniciu, "An algorithm for data-driven bandwidth selection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 2, pp. 281–288, 2003.

[10] D. Comaniciu and P. Meer, "Mean shift a robust approach toward feature space analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 5, pp. 603–619, 2002.

[11] M. Piccardi and T. Jan, "Efficient mean-shift backgonmd subtraction," *Proc. of IEEE KIP*, Oct. 2004.

[12] B. Han, D. Comaniciu, and L. S. Davis, "Sequential kernel density approximation through mode propagation: applications to background modeling," *Proceedings of the Asian Conference on Computer Vision*, Jan. 2004.

[13] M. Seki, T. Wada, H. Fujiwara, and K. Sumi, "Background subtraction based on cooccurrence of image variations," *Proceedings of the International Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 65–72, 2003.

[14] N. M. Oliver, B. Rosario, and A. P. Pentland, "A bayesian computer vision system for modeling human interactions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 831–843, 2000.

[15] K. V. Mardia, J. T. Kent, and J. M. Bibby, *Multivariate Analysis*. Academic Press, London, 1979.

[16] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, pp. 2323–2326, December 2000.

[17] J. B. Tenenbaum, V. de Silva, and J. C. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, pp. 2319–2323, December 2000.

[18] R. R. Coifman and S. Lafon, "Diffusion maps," *Applied and Computational Harmonic Analysis*, vol. 21, pp. 5–30, 2006.

[19] F. R. K. Chung, *Spectral Graph Theory*. AMS Regional Conference Series in Mathematics, 92, 1997.

[20] K. Toyama, J. K. B. Brumitt, and B. Meyers, "Wallflower: Principles and practice of background maintenance," *International Conference on Computer Vision*, pp. 255–261, 1999.

[21] D. Koller, J. W. Weber, and J. Malik, "Robust multiple car tracking with occlusion reasoning," *The European Conference on Computer Vision*, 1994.